

Efficient State Management in Distributed Ledgers

Dimitris Karakostas

Nikos Karayiannidis

Aggelos Kiayias

Types of shared state

1. *Public ledger*: the ordered list of published transactions
2. *Mempool*: the set of unpublished transactions
3. **Active state**:
 - a. UTxO set
 - b. Global state (Ethereum)

Attacks against active state



The Ethereum network is currently undergoing a DoS attack

Posted by Jeffrey Wilcke on September 22, 2016

Research & Development

July 2015 flood attack

The **July 2015 flood attack** was a large "stress test" of the Bitcoin network. The possibility was subsequent to stress tests executed in June.

Overview

- UTxO-based ledger model
- Transaction-oriented state optimization
 - Rule based (heuristic)
 - Transformations
 - Exhaustive search of transaction space
- Fee model
 - The *state efficiency* property
 - A state efficient Bitcoin

UTxO Ledger

- $\mathcal{L} \stackrel{\text{def}}{=} \text{List}[\text{Transaction}]$
- $\text{Transaction} \stackrel{\text{def}}{=} (\text{inputs} : \text{Set}[\text{Input}],$
 $\text{outputs} : \text{List}[\text{Output}],$
 $\text{forge} : \text{Value},$
 $\text{fee} : \text{Value})$
- $\text{UTxO} \stackrel{\text{def}}{=} (\alpha : \text{Address}, \text{value} : \text{Value}, \text{created} : \text{Timestamp})$
- $\text{Input} \stackrel{\text{def}}{=} (\text{id} : \text{Hash}, \text{index} : \text{Int})$
- $\text{State} \stackrel{\text{def}}{=} \text{Set}[\text{Input}]$
- $\text{costTx} : \text{UtxoTx} \rightarrow \text{LedgerState} \rightarrow \text{Cost}$
 $\text{costTx}(\tau, \mathcal{S}) = \text{cost}(\text{txRun}(\tau, \mathcal{S})) - \text{cost}(\mathcal{S})$

Transaction Model

- Transaction equivalency:

$$\forall \alpha \in A \quad \sum_{\substack{i \in \mathcal{S}_1 \\ o_i = \text{out}(i, \mathcal{L}) \\ o_i.\text{address} = \alpha}} o_i.\text{value} = \sum_{\substack{j \in \mathcal{S}_2 \\ o_j = \text{out}(j, \mathcal{L}) \\ o_j.\text{address} = \alpha}} o_j.\text{value}$$

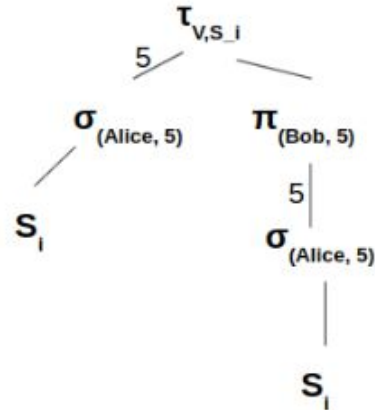
where A is the set of all addresses of the parties participating in the distributed ledger system.

- **Input Selection** $\sigma_{[(a_1, v_1), \dots, (a_n, v_n)]} : \text{LedgerState} \rightarrow \text{LedgerState}$
- **Output Creation** $\pi_{[(a_1, v_1), \dots, (a_n, v_n)]} : \text{LedgerState} \rightarrow \text{LedgerState}$
- **Transaction Validation** $\tau_{V, S_i} : \text{LedgerState} \rightarrow \text{LedgerState} \rightarrow \text{LedgerState}$

Alice wants to give Bob 5 coins

- Alice: {50, 15, 5, 2, 2, 1} (75 total coins, 6 UTXOs)
- Bob: {4, 3, 3, 1, 1} (12 total coins, 5 UTXOs)


1. Alice -5-> Bob



$S_i =$ Alice {50,15,2,2,1} 70(5)
Bob {5,4,3,3,1,1} 17(6)
TotVal = 87, **Cost = 11**

Alice wants to give Bob 5 coins


- Alice: {50, 15, 5, 2, 2, 1} (75 total coins, 6 UTXOs)
- Bob: {4, 3, 3, 1, 1} (12 total coins, 5 UTXOs)



Alice wants to give Bob 5 coins

- Alice: {50, 15, 5, 2, 2, 1} (75 total coins, 6 UTXOs)
- Bob: {4, 3, 3, 1, 1} (12 total coins, 5 UTXOs)

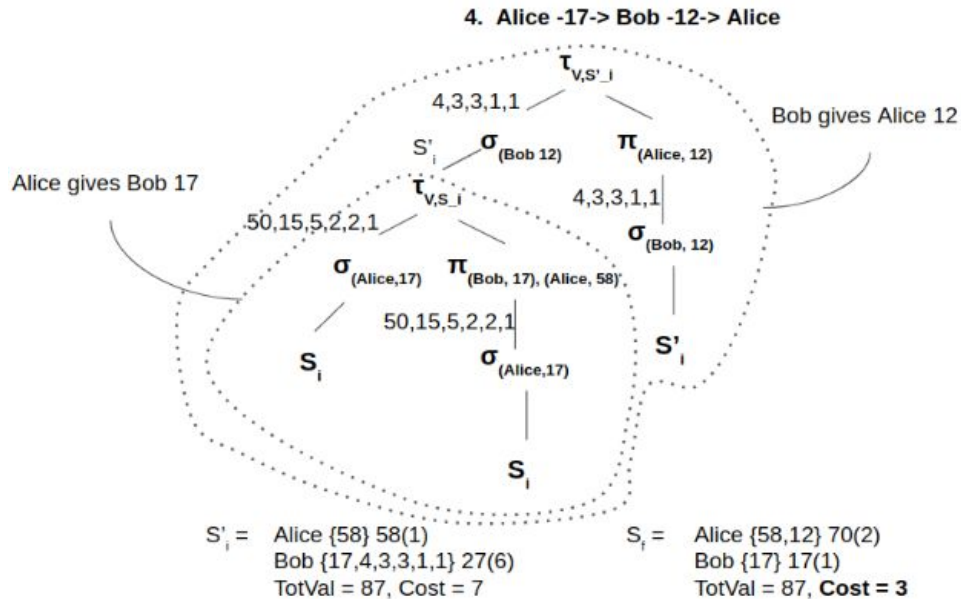
3. Alice -2,2,1-> Bob



$S_i =$ Alice {50,15,5} 70(3)
Bob {5,4,3,3,1,1} 17(6)
TotVal = 87, Cost = 9

Alice wants to give Bob 5 coins

- Alice: {50, 15, 5, 2, 2, 1} (75 total coins, 6 UTxOs)
- Bob: {4, 3, 3, 1, 1} (12 total coins, 5 UTxOs)



Transaction Optimization

- Rules
 - One output per address
 - Consume maximum # of inputs
 - ...
- Transformations
 - Reorder tx creation to minimize cost
 - ...
- Methods
 - Input selection algorithm
 - ...
- Search of transaction space
 - Optimization metrics (eg. cost, size)
 - Limits (eg. max fees to spend)

State-aware input selection

- *Problem*: select inputs to consume
- *Goal*: minimize state cost + upper cap on fees
- *Fee model*:
 - input & output sizes fixed
 - $\text{fees} = (\text{numOf}(\text{inputs}) + \text{numOf}(\text{outputs}) + 1) * \text{utxoSizeBytes} * \text{costPerByte}$
- *State-aware input selection algorithm*:
 - Sort inputs per value
 - Find largest window such that:
 - total value of inputs \geq tx output
 - fees of consuming inputs \leq upper cap

State Efficient Fee Model

State Efficiency:

$$\forall \mathcal{S} \in \mathcal{S} \forall \tau_1, \tau_2 \in \mathbb{T} \mid \tau_1 \equiv \tau_2 \wedge \text{costTx}(\tau_1, \mathcal{S}) > \text{costTx}(\tau_2, \mathcal{S}) : F(\tau_1, \mathcal{S}) > F(\tau_2, \mathcal{S})$$

Bitcoin

- $F(\text{tx}) = \text{sizeof}(\text{tx}) * \text{feePerByte}$
- Simple model:
 - i : size of consuming an input
 - o : size of UTxO
- Example:
 - tx_1 : 1 input, 1 output
 - $\text{cost}(\text{tx}_1) = 0$
 - $F(\text{tx}_1) = (i + o) * \text{feePerByte}$
 - tx_2 : 2 inputs, 1 output
 - $\text{cost}(\text{tx}_2) = 1$
 - $F(\text{tx}_2) = (2i + o) * \text{feePerByte} = F(\text{tx}_1) + i * \text{feePerByte}$

State-efficient Bitcoin

- $F'(tx) = \text{sizeof}(tx) * \text{feePerByte} + \text{cost}(tx) * \text{feePerUTxO}$
- Intuition:
 - When a user creates a UTxO (increasing the state), they pay an extra fee
 - When a user consumes an input (decreasing the state), they get reimbursed
- If $\text{feePerUTxO} > i * \text{feePerByte}$, fee function is state efficient
 - The user who creates a UTxO reimburses the (future) user who consumes it for at least enough fees to cover the increase in tx size

Ethereum

- **selfdestruct:**
 - destroys the smart contract
 - frees up global state
- Programmer who writes the *selfdestruct* line of code pays fee for it (due to increased smart contract size)
- Gives user who *runs* it with (gas) fee discount
- User who receives discount \neq User who pays the fee
- Discount is *at most* half transaction's fees
 - User has no incentive to *only* run *selfdestruct*

Summary

- Poor state management may lead to:
 - Bad performance, i.e. bloated state shared across all parties
 - Denial-of-Service
- Transaction optimization should take into account state efficiency
- Fees should reflect state cost of transaction
- Paper's full version available: <https://eprint.iacr.org/2020/525.pdf>

Thank you!